Math 562 Fall 2020

Renming Song

University of Illinois at Urbana-Champaign

September 25, 2020

Outline

Outline

HW3 is due on Friday, 10/02, at noon. Please submit your HW3 via the course Moodle page. Make sure that your HW is uploaded successfully.

Slides and videos of the lectures are also available from the Moodle page.

HW3 is due on Friday, 10/02, at noon. Please submit your HW3 via the course Moodle page. Make sure that your HW is uploaded successfully.

Slides and videos of the lectures are also available from the Moodle page.

Outline

Applications

(d) For every $a > 0$, consider the stopping time $U_a = \inf\{t \geq 0 : |B_t| = a\}$. Then for every $\lambda > 0$,

$$\mathbb{E}[\exp(-\lambda U_a)] = \frac{1}{\cosh(a\sqrt{2\lambda})}.$$

First note that U_a and B_{U_a} are independent, since by symmetry

$$\mathbb{E}[1_{\{B_{U_a}=a\}} e^{-\lambda U_a}] = \mathbb{E}[1_{\{B_{U_a}=-a\}} e^{-\lambda U_a}] = \frac{1}{2}\mathbb{E}[e^{-\lambda U_a}].$$

Again, we consider the exponential martingale

$$N_t^\lambda = \exp\left(\lambda B_t - \frac{\lambda^2}{2} t\right).$$

Applications

(d) For every $a > 0$, consider the stopping time $U_a = \inf\{t \geq 0 : |B_t| = a\}$. Then for every $\lambda > 0$,

$$\mathbb{E}[\exp(-\lambda U_a)] = \frac{1}{\cosh(a\sqrt{2\lambda})}.$$

First note that U_a and B_{U_a} are independent, since by symmetry

$$\mathbb{E}[1_{\{B_{U_a}=a\}} e^{-\lambda U_a}] = \mathbb{E}[1_{\{B_{U_a}=-a\}} e^{-\lambda U_a}] = \frac{1}{2}\mathbb{E}[e^{-\lambda U_a}].$$

Again, we consider the exponential martingale

$$N_t^\lambda = \exp\left(\lambda B_t - \frac{\lambda^2}{2} t\right).$$

Noe that $N^{\lambda}_{t\wedge U_a}$ is bounded above $e^{\lambda a}$, so by the optional stopping theorem,

$$\cosh(a\lambda)\mathbb{E}[e^{-\frac{\lambda^2}{2}U_a}] = \mathbb{E}[N^{\lambda}_{U_a}] = \mathbb{E}[N^{\lambda}_0] = 1.$$

Replacing λ by $\sqrt{2\lambda}$, we immediately get the desired assertion.

We end this chapter with the optional stopping theorem for non-negative supermartingales. This result will be useful in later applications to Markov processes. We first note that, if $(Z_t)_{t\geq 0}$ is a non-negative supermartingale with right-continuous sample paths, $(Z_t)_{t\geq 0}$ is automatically bounded in L^1 since $\mathbb{E}[Z_t] \leq \mathbb{E}[Z_0]$, and by Theorem 3.19, Z_t converges a.s. to a random variable Z_∞ as $t \to \infty$. Thus we can define Z_T for any stopping time T.

Noe that $N^\lambda_{t\wedge U_a}$ is bounded above $e^{\lambda a}$, so by the optional stopping theorem,

$$\cosh(a\lambda)\mathbb{E}[e^{-\frac{\lambda^2}{2}U_a}] = \mathbb{E}[N^\lambda_{U_a}] = \mathbb{E}[N^\lambda_0] = 1.$$

Replacing λ by $\sqrt{2\lambda}$, we immediately get the desired assertion.

We end this chapter with the optional stopping theorem for non-negative supermartingales. This result will be useful in later applications to Markov processes. We first note that, if $(Z_t)_{t\geq 0}$ is a non-negative supermartingale with right-continuous sample paths, $(Z_t)_{t\geq 0}$ is automatically bounded in L^1 since $\mathbb{E}[Z_t] \leq \mathbb{E}[Z_0]$, and by Theorem 3.19, Z_t converges a.s. to a random variable Z_∞ as $t \to \infty$. Thus we can define Z_T for any stopping time T.

Theorem 3.25

Let $(Z_t)_{t\geq 0}$ be a non-negative supermartingale with right-continuous sample paths. Let U and V be two stopping times such that $U \leq V$. Then, Z_U and Z_V are in L^1, and

$$Z_U \geq \mathbb{E}[Z_V|\mathcal{F}_U].$$

Remark

This implies that $\mathbb{E}[Z_U] \geq \mathbb{E}[Z_V]$, and since $Z_U = Z_V = Z_\infty$ on $\{U = \infty\}$, we also have

$$\mathbb{E}[1_{\{U<\infty\}}Z_U] \geq \mathbb{E}[1_{\{U<\infty\}}Z_V] \geq \mathbb{E}[1_{\{V<\infty\}}Z_V].$$

Theorem 3.25

Let $(Z_t)_{t\geq 0}$ be a non-negative supermartingale with right-continuous sample paths. Let U and V be two stopping times such that $U \leq V$. Then, Z_U and Z_V are in L^1, and

$$Z_U \geq \mathbb{E}[Z_V|\mathcal{F}_U].$$

Remark

This implies that $\mathbb{E}[Z_U] \geq \mathbb{E}[Z_V]$, and since $Z_U = Z_V = Z_\infty$ on $\{U = \infty\}$, we also have

$$\mathbb{E}[1_{\{U<\infty\}}Z_U] \geq \mathbb{E}[1_{\{U<\infty\}}Z_V] \geq \mathbb{E}[1_{\{V<\infty\}}Z_V].$$

Proof of Theorem 3.25

Step 1: We will show that, under the extra assumption U and V are bounded, $\mathbb{E}[Z_U] \geq \mathbb{E}[Z_V]$. Let $p \geq 1$ be an integer such that $V \leq p$. For every integer $n \geq 0$, define

$$U_n = \sum_{k=0}^{p2^n-1} \frac{k+1}{2^n} 1_{\{k2^{-n} < U \leq (k+1)2^{-n}\}}$$

and

$$V_n = \sum_{k=0}^{p2^n-1} \frac{k+1}{2^n} 1_{\{k2^{-n} < V \leq (k+1)2^{-n}\}}.$$

Then (U_n) and (V_n) are two sequences of bounded stopping times that decrease respectively to U and V, and $U_n \leq V_n$ for each $n \geq 0$. The right-continuity of sample paths ensures that $Z_{U_n} \to Z_U$ and $Z_{V_n} \to Z_V$ a.s. as $n \to \infty$.

Proof of Theorem 3.25 (cont)

By the optional stopping theorem for discrete supermartingales in the case of bounded stopping times, with respect to the filtration $(\mathcal{F}_{k2^{-n-1}})_{k\geq 0}$, we have for every

$$Z_{U_{n+1}} \geq \mathbb{E}[Z_{U_n}|\mathcal{F}_{U_{n+1}}].$$

Let $Y_n = Z_{U_{-n}}$ and $\mathcal{H}_n = \mathcal{F}_{U_{-n}}$ for every $n \geq 0$. Then $(Y_n)_{n\geq 0}$ is a backward supermartingale with respect to the filtration $(\mathcal{H}_n)_{n\leq 0}$. Since, by the optional stopping theorem, $\mathbb{E}[Z_{U_n}] \leq \mathbb{E}[Z_0]$ for every $n \geq 0$, the sequence $(Y_n)_{n\geq 0}$ is bounded in L^1, and by the convergence theorem for backward supermartingales, it converges in L^1. Hence the convergence of Z_{U_n} to Z_U also holds in L^1 and similarly Z_{V_n} converges to Z_V in L^1. Since $U_n \leq V_n$, using the optional stopping theorem again, we get $\mathbb{E}[Z_{U_n}] \geq \mathbb{E}[Z_{V_n}]$. Letting $n \to \infty$, we get $\mathbb{E}[Z_U] \geq \mathbb{E}[Z_V]$.

Proof of Theorem 3.25 (cont)

Now we use to Step 1 to finish the proof of the theorem. By Step 1 applied to the stopping times 0 and $U \wedge p$, we have $\mathbb{E}[Z_{U\wedge p}] \le \mathbb{E}[Z_0]$ for every $p \ge 1$, and Fatou's lemma gives $\mathbb{E}[Z_U] \le \mathbb{E}[Z_0] < \infty$ and similarly $\mathbb{E}[Z_V] < \infty$. Fix $A \in \mathcal{F}_U$ and recall $U^A := U 1_A + \infty \cdot 1_{A^c}$ is a stopping time. By Step 1, we have, for every $p \ge 1$,

$$\mathbb{E}[Z_{U^A \wedge p}] \ge \mathbb{E}[Z_{V^A \wedge p}].$$

By writing each of these two expectations as a sum of expectations over the sets A^c, $A \cap \{U \le p\}$ and $A \cap \{U > p\}$, and noting that $U > p$ implies $V > p$, we get

$$\mathbb{E}[Z_U 1_{A\cap\{U\le p\}}] \ge \mathbb{E}[Z_V 1_{A\cap\{U\le p\}}].$$

Proof of Theorem 3.25 (cont)

Letting $p \uparrow \infty$, we get

$$\mathbb{E}[Z_U 1_{A\cap\{U<\infty\}}] \geq \mathbb{E}[Z_V 1_{A\cap\{U<\infty\}}].$$

On the other hand, the equality

$$\mathbb{E}[Z_U 1_{A\cap\{U=\infty\}}] = \mathbb{E}[Z_V 1_{A\cap\{U=\infty\}}]$$

is trivial and by adding it to the preceding display, we obtain

$$\mathbb{E}[Z_U 1_A] \geq \mathbb{E}[Z_V 1_A].$$

Since this holds for every $A \in \mathcal{F}_U$ and Z_U is $\mathcal{F}_U$-measurable, the desired result follows.

Outline

In this chapter, we study continuous semimartingales. By definition, a continuous semimartingale is the sum of a continuous local martingale and a (continuous) finite variation process. We will study separately these two classes of processes. We start with some preliminaries about deterministic functions with finite variation, before considering the corresponding processes. We then define (continuous) local martingales and we construct the quadratic variation of a local martingale, which will play a fundamental role in the construction of stochastic integrals

We explain how properties of a local martingale are related to those of its quadratic variation. Finally, we introduce continuous semimartingales and their quadratic variation processes.

In this chapter, we study continuous semimartingales. By definition, a continuous semimartingale is the sum of a continuous local martingale and a (continuous) finite variation process. We will study separately these two classes of processes. We start with some preliminaries about deterministic functions with finite variation, before considering the corresponding processes. We then define (continuous) local martingales and we construct the quadratic variation of a local martingale, which will play a fundamental role in the construction of stochastic integrals

We explain how properties of a local martingale are related to those of its quadratic variation. Finally, we introduce continuous semimartingales and their quadratic variation processes.

Since we are only interested in continuous semimartinagles, we restrict our attention to continuous functions of finite variation. Recall that a signed measure on a compact interval $[0, T]$ is the difference of two finite positive measures on $[0, T]$.

Definition 4.1

Let $T > 0$. A continuous function $a : [0, T] \to \mathbb{R}$ with $a(0) = 0$ is said to be of finite variation if there exists a signed measure μ on $[0, T]$ such that $a(t) = \mu([0, t])$ for every $t \in [0, T]$.

The measure μ is then determined uniquely by a. Since a is continuous and $a(0) = 0$, μ has no atoms.

Since we are only interested in continuous semimartinagles, we restrict our attention to continuous functions of finite variation. Recall that a signed measure on a compact interval $[0, T]$ is the difference of two finite positive measures on $[0, T]$.

Definition 4.1

Let $T > 0$. A continuous function $a : [0, T] \to \mathbb{R}$ with $a(0) = 0$ is said to be of finite variation if there exists a signed measure μ on $[0, T]$ such that $a(t) = \mu([0, t])$ for every $t \in [0, T]$.

The measure μ is then determined uniquely by a. Since a is continuous and $a(0) = 0$, μ has no atoms.

Since we are only interested in continuous semimartinagles, we restrict our attention to continuous functions of finite variation. Recall that a signed measure on a compact interval $[0, T]$ is the difference of two finite positive measures on $[0, T]$.

Definition 4.1

Let $T > 0$. A continuous function $a : [0, T] \to \mathbb{R}$ with $a(0) = 0$ is said to be of finite variation if there exists a signed measure μ on $[0, T]$ such that $a(t) = \mu([0, t])$ for every $t \in [0, T]$.

The measure μ is then determined uniquely by a. Since a is continuous and $a(0) = 0$, μ has no atoms.

The general definition of a function with finite variation does not require continuity nor the condition $a(0) = 0$. We impose these two conditions for convenience.

The decomposition of μ as a difference of two finite positive measures on $[0, T]$ is not unique, but there exists a unique decomposition $\mu = \mu_+ - \mu_-$ such that μ_+ and μ_- are supported on disjoint Borel sets. To get the existence of such a decomposition, start from an arbitrary decomposition $\mu = \mu_1 - \mu_2$, define $\nu = \mu_1 + \mu_2$ and then use the Radon-Nikodym theorem to find two non-negative Borel functions h_1 and h_2 on $[0, T]$ such that

$$\mu_1(dt) = h_1(t)\nu(dt), \quad \mu_2 d(t) = h_2(t)\nu(dt).$$

The general definition of a function with finite variation does not require continuity nor the condition $a(0) = 0$. We impose these two conditions for convenience.

The decomposition of μ as a difference of two finite positive measures on $[0, T]$ is not unique, but there exists a unique decomposition $\mu = \mu_+ - \mu_-$ such that μ_+ and μ_- are supported on disjoint Borel sets. To get the existence of such a decomposition, start from an arbitrary decomposition $\mu = \mu_1 - \mu_2$, define $\nu = \mu_1 + \mu_2$ and then use the Radon-Nikodym theorem to find two non-negative Borel functions h_1 and h_2 on $[0, T]$ such that

$$\mu_1(dt) = h_1(t)\nu(dt), \quad \mu_2 d(t) = h_2(t)\nu(dt).$$

Then, if $h(t) = h_1(t) - h_2(t)$, we have

$$\mu(dt) = h(t)\nu(dt) = h^+(t)\nu(dt) - h^-(t)\nu(dt),$$

which gives the decomposition $\mu = \mu_+ - \mu_-$ with $\mu_+(dt) = h^+(t)\nu(dt)$ and $\mu_-(dt) = h^-(t)\nu(dt)$, and the measures μ_+ and μ_- are respectively supported on the disjoint Borel sets $D_+ := \{t : h(t) > 0\}$ and $D_- := \{t : h(t) < 0\}$. The uniqueness of this decomposition $\mu = \mu_+ - \mu_-$ follows from the fact that we have necessarily, for every $A \in \mathcal{B}([0, T])$,

$$\mu_+(A) = \sup\{\mu(C) : C \in \mathcal{B}([0, T]), C \subset A\}.$$

We write $|\mu|$ for the (finite) positive measure $|\mu| = \mu_+ + \mu_-$. The measure $|\mu|$ is called the total variation of a. We have $|\mu(A)| \leq |\mu|(A)$ for every $A \in \mathcal{B}([0, T])$.

Then, if $h(t) = h_1(t) - h_2(t)$, we have

$$\mu(dt) = h(t)\nu(dt) = h^+(t)\nu(dt) - h^-(t)\nu(dt),$$

which gives the decomposition $\mu = \mu_+ - \mu_-$ with $\mu_+(dt) = h^+(t)\nu(dt)$ and $\mu_-(dt) = h^-(t)\nu(dt)$, and the measures μ_+ and μ_- are respectively supported on the disjoint Borel sets $D_+ := \{t : h(t) > 0\}$ and $D_- := \{t : h(t) < 0\}$. The uniqueness of this decomposition $\mu = \mu_+ - \mu_-$ follows from the fact that we have necessarily, for every $A \in \mathcal{B}([0, T])$,

$$\mu_+(A) = \sup\{\mu(C) : C \in \mathcal{B}([0, T]), C \subset A\}.$$

We write $|\mu|$ for the (finite) positive measure $|\mu| = \mu_+ + \mu_-$. The measure $|\mu|$ is called the total variation of a. We have $|\mu(A)| \leq |\mu|(A)$ for every $A \in \mathcal{B}([0, T])$.

Moreover, the Radon-Nikodym derivative of μ with respect to $|\mu|$ is

$$\frac{d\mu}{d|\mu|} = 1_{D_+} - 1_{D_-}.$$

The fact that $a(t) = \mu_+([0,t]) - \mu_-([0,t])$ shows that a is the difference of two monotone non-decreasing continuous functions that vanish at 0 (since μ has no atoms, the same holds for μ_+ and μ_-). Conversely, the difference of two monotone nondecreasing continuous functions that vanish at 0 has finite variation in the sense of the previous definition. Indeed, this follows from the well-known fact that the formula $g(t) = \theta([0,t]), t \in [0,T]$ induces a bijection between monotone nondecreasing right-continuous functions $g : [0,T] \to \mathbb{R}_+$ and finite positive measures θ on $[0,T]$.

Moreover, the Radon-Nikodym derivative of μ with respect to $|\mu|$ is

$$\frac{d\mu}{d|\mu|} = 1_{D_+} - 1_{D_-}.$$

The fact that $a(t) = \mu_+([0,t]) - \mu_-([0,t])$ shows that a is the difference of two monotone non-decreasing continuous functions that vanish at 0 (since μ has no atoms, the same holds for μ_+ and μ_-). Conversely, the difference of two monotone nondecreasing continuous functions that vanish at 0 has finite variation in the sense of the previous definition. Indeed, this follows from the well-known fact that the formula $g(t) = \theta([0,t]), t \in [0,T]$ induces a bijection between monotone nondecreasing right-continuous functions $g : [0,T] \to \mathbb{R}_+$ and finite positive measures θ on $[0,T]$.

Let $f : [0, T] \to \mathbb{R}$ be a measurable function such that $\int_{[0,T]} |f(s)||\mu|(ds < \infty$. We define

$$\int_0^T f(s)da(s) = \int_{[0,T]} f(s)\mu(ds)$$
$$\int_0^T f(s)|da(s)| = \int_{[0,T]} f(s)|\mu|(ds).$$

Then the bound

$$|\int_0^T f(s)da(s)| \leq \int_0^T |f(s)||da(s)|$$

holds. By restricting a to $[0, t]$, we can define $\int_0^t f(s)da(s)$ for every $t \in [0, T]$, and we note that the function $t \mapsto \int_0^t f(s)da(s)$ is of finite variation on $[0, T]$.

Proposition 4.2

For every $t \in [0, T]$,

$$\int_0^t |da(s)| = \sup\{\sum_{j=1}^{p} |f(t_j) - f(t_{j-1})|\},$$

where the supremum is over all partitions $0 = t_0 < t_1 < \cdots < t_p = t$ of $[0, t]$. More precisely, for any increasing sequence $0 = t_0^n < t_1^n < \cdots < t_{p_n}^n = t$ of partitions of $[0, t]$ whose mesh tends to 0, we have

$$\lim_{n\to\infty} \sum_{j=1}^{p_n} |f(t_j^n) - f(t_{j-1}^n)| = \int_0^t |da(s)|.$$

In the usual presentation of functions with finite variation, one starts from the property that the supremum in the first display of the proposition is finite.

Proposition 4.2

For every $t \in [0, T]$,

$$\int_0^t |da(s)| = \sup\{\sum_{j=1}^{p} |f(t_j) - f(t_{j-1})|\},$$

where the supremum is over all partitions $0 = t_0 < t_1 < \cdots < t_p = t$ of $[0, t]$. More precisely, for any increasing sequence $0 = t_0^n < t_1^n < \cdots < t_{p_n}^n = t$ of partitions of $[0, t]$ whose mesh tends to 0, we have

$$\lim_{n\to\infty} \sum_{j=1}^{p_n} |f(t_j^n) - f(t_{j-1}^n)| = \int_0^t |da(s)|.$$

In the usual presentation of functions with finite variation, one starts from the property that the supremum in the first display of the proposition is finite.